

Effects of New York City’s Neighborhood Policing Policy

August 2021 pre-print version

For the copy-edited, formatted version in *Police Quarterly* see:

<https://doi.org/10.1177/10986111211046991>

Brenden Beck¹, Joseph Antonelli², Gabriela Piñeros³

Abstract

Between 2015 and 2018, New York City adopted “neighborhood policing,” an expansive policy to encourage interactions between police officers and community members. Among other changes, the initiative established hundreds of new “neighborhood-coordination” officers and gave “steady-sector” officers time away from 911 response to dedicate to resident interactions. This study evaluates the initiative’s effects on crime, complaints of police misconduct, racial disparities, and arrests. Using monthly data on New York City’s 76 police precincts between 2006 and 2019, we estimate the policy’s causal effect using high-dimensional time series models. This approach accounts for the policy’s staggered adoption, addresses potential correlation among outcomes and between precincts, and controls for unobserved precinct characteristics. We find neighborhood policing reduced misdemeanor and proactive arrests, especially in higher-poverty precincts outside of Manhattan, though it did not change the racial disparities of such arrests. The policy did not affect crime. It briefly increased complaints against police.

Keywords community engagement; misdemeanor arrests; staggered policy adoption; heterogeneous treatment effects

Acknowledgements We thank Rebecca Hanson for her insightful comments on an earlier draft, and we thank Kaitlyn Palm and Abigail Grzyb for their research assistance.

¹ Corresponding Author: Brenden.Beck@ucdenver.edu. University of Colorado Denver, Department of Sociology, Campus Box 105, PO Box 173364, Denver, CO 80217.

² University of Florida, Department of Statistics

³ University of Florida, Department of Sociology, Criminology, and Law

Underlying code and data are available at:

<https://www.dropbox.com/sh/loeaacy93zpf6no/AABpAPRueVFL8AHLFQCX-i9ea>

Introduction

In 2015, the New York Police Department (NYPD) began a new program called “neighborhood policing” aimed at reducing crime and improving police-community relations. At the launch event for the policy, Mayor Bill de Blasio said the “truly transformative” program was

being applied “to ensure the people of this city have a police force that is deeply connected at the neighborhood level...to build relationships and deepen trust” (NYPD 2015b). The City’s police commissioner and an architect of the policy, Bill Bratton, characterized neighborhood policing as “a new era” (NYPD 2015b). By 2018, the initiative had restructured each of the NYPD’s precincts into smaller sectors; it appointed hundreds of new “neighborhood-coordination” officers whose training and primary duties were in community engagement; and it established “steady-sector” officers who walked a consistent beat and had a third of their shift away from emergency response to devote to problem-solving work and resident interaction (NYPD 2018). The policy is spreading beyond New York. Chicago announced in 2020 that it was adopting a neighborhood policing initiative modeled on New York City’s (Office of the Mayor of Chicago 2020).

In this study, we evaluate neighborhood policing’s impacts on reported crime, civilian complaints of police misconduct, low-level arrests, and racial disparities in low-level arrests. The study is, to our knowledge, the first systematic assessment of New York City’s neighborhood policing program. To perform the analysis, we gather monthly data on New York City’s 76 police precincts from 2006 to 2019 and estimate high-dimensional time series models. We validate our approach using the data from before the initiation of neighborhood policing. We propose policy recommendations and address methodological challenges in evaluating initiatives that, like this one, have precincts adopt the policy at different times, have correlated outcomes, and have effects that vary by unit-level characteristics (treatment effect heterogeneity).

New York’s neighborhood policing both draws on and departs from past community policing initiatives (Reisig 2010; Bratton 2015; Oliver 2000). The wide variety of interventions that have been labeled “community policing” has sometimes muddled the concept (Connell, Miggins, & McGloin 2008), but there are common elements the “ideal” program should have:

“community partnerships, organizational transformation, and problem solving” (Gill et al., 2014, 410; Weisburd & Eck 2004; Cordner, 1997). In this respect, New York’s version is typical. It promotes officer liaisons with civilians, it reorganized the city’s precincts, and it appointed officers who specialize in problem solving (Bratton 2015; NYPD 2018). Several elements, however, distinguish New York’s neighborhood policing from many previous community policing initiatives. First, it enlists specialized teams of “neighborhood-coordination” and “steady-sector” officers as well as traditional patrol officers in community engagement. This eliminates the typical “bifurcation” between community police who solve problems and patrol police who answer calls for service (Bratton 2015, 2; Gill et al. 2014). A second distinguishing feature is neighborhood policing’s scale. It was not deployed only to crime hot-spots, only in communities with strained police-community relations, or only in sectors with willing officers, as some interventions have been (Weisburd et al. 2020; Connell, Miggans, & McGloin 2008), but was used throughout the city and with all officers (NYPD 2018). The program coincided with the hiring of 1,300 additional officers and much of this increased capacity was devoted to neighborhood policing (NYPD 2015a). This represents an increase larger than most cities’ police forces (Hyland & Davis 2019).

New York City has often been a bellwether department, setting policy trends for many other departments, and it might again with neighborhood policing. After the city adopted broken windows policing in the 1990s (a policy also designed by Bill Bratton), the misdemeanor-focused approach spread to many other cities (Lum & Nagin 2017). Other cities are likely to follow Chicago in adopting neighborhood policing, so evidence of the policy’s efficacy is vital. Yet, policymakers considering whether to adopt the policy currently lack evidence to guide their decision. As of August 2021, New York City had not yet released its own promised evaluation of neighborhood policing, and it cancelled a public opinion poll assessing the policy’s efficacy

(Edelman, 2020). New York City's size and the variety of its neighborhoods provide a large sample of communities with a wide range of demographics and crime levels, allowing us to observe how the policy's effects vary across diverse places.

Neighborhood policing's architects hoped that by improving cooperation with community members, they would solve and deter more crimes. They made reducing crime the "first goal" of neighborhood policing (NYPD, 2018, 4), so crime is the first outcome we measure. We also suspect the policy will have consequences for low-level arrests. Neighborhood policing increases police contact with community members, and this might escalate discretionary arrests because it increases opportunities for arrest and it overlays—rather than replaces—broken windows policing, a strategy involving high levels of misdemeanor arrests (Gascón & Rousell 2019). Understanding the policy's impact on arrests is especially important in light of the growing body of research demonstrating the serious consequences of misdemeanor arrests for people arrested, their communities, and police (Natapoff 2018). These consequences usually fall disproportionately on non-white people (Committee on Proactive Policing 2018), so we analyze neighborhood policing's impacts on the frequency and the racial equity of low-level arrests.

Finally, we want to capture whether neighborhood policing affects civilian complaints of police misconduct, and hypothesize they, like low-level arrests, will increase as contacts between police and residents increase. Neighborhood policing encourages police to use their discretion to solve problems. The increase in autonomy and community interactions that neighborhood policing encourages can introduce more opportunities for excessive force or discriminatory policing, problems that can generate complaints (Hickman, Piquero, & Greene 2000; Kessler 1999).

In choosing an analytic approach to evaluate neighborhood policing's impact on these four outcomes, we want to use a method that relies on as few assumptions as possible and allows its

results to be assessed for biases. In this case, we also have to account for the policy's adoption by different precincts at different times and for any impacts that vary by precinct characteristics. In a context like this, researchers have typically used difference in differences (Ashenfelter 1978), synthetic controls (Abadie 2010), or interrupted time series models (Lopez Bernal, Cummins, & Gasparrini 2017). While these approaches have been effective in other cases, none are a perfect fit for the situation at hand. Difference-in-difference designs have only recently been developed for situations with staggered policy adoption, and such approaches do not assess treatment effect heterogeneity. Synthetic control approaches require at least some units to never implement the policy, but in New York City every precinct eventually adopted neighborhood policing. Interrupted time series designs offer the most promise, though they still rely on the assumption that there are no unmeasured time-varying covariates that are associated with both the treatment status of a precinct and the outcome of interest. Further, none of these approaches account for potential correlation of outcomes across units. Ignoring such correlation could lead to overly small confidence intervals and misleading statements of statistical significance.

To address these methodological challenges, we apply a newly developed approach (Antonelli & Beck 2021) that simultaneously accounts for all these concerns. The approach uses Bayesian time series models to predict what would have happened to each New York City precinct in the absence of neighborhood policing. This allows us to estimate the difference between the outcomes we observe after policy initiation and the unknown counterfactual values that would have occurred in the absence of the policy. This time-series approach accounts for both temporal correlation across time as well as spatial correlation between precincts. Further, it allows us to use flexible modeling techniques to reveal how the effects of neighborhood policing vary by precinct characteristics such as socioeconomic status and racial composition. Crucially, the approach is

robust to biases arising from unmeasured covariates that affect the treatment and the outcome, frequently referred to as confounders. As a sensitivity analysis, we compare our findings to those from interrupted time series models to check against an approach with different assumptions, and found similar results, increasing our confidence in our results.

Our primary time series approach relies on a stationarity assumption, which states that the estimated model continues to hold into time periods after neighborhood policing is adopted. Unlike assumptions about the presence of unmeasured confounders, we can assess the plausibility of this assumption. Using a simulation study based solely on the observed data in the time period before policy initiation, we can assess whether the stationarity assumption holds in the time periods leading up to treatment initiation. While this does not guarantee that our assumption holds, it provides increased confidence that our approach and modeling assumptions are reasonable in the data set being used. The simulation indicates our proposed approach provides accurate estimates and generates confidence intervals that control type I error rates.

Neighborhood policing was implemented against the backdrop of several shifts in policing, both locally and nationally. In the 2000s and early 2010s, New York City aggressively employed a strategy called “stop and frisk” whereby police made frequent searches of pedestrians to look for illegal guns and drugs. Under pressure from social movements and a federal court order, the NYPD sharply curtailed stop and frisk starting in 2012, reducing the number of stops from a high of 680,000 in 2011 to fewer than 46,000 in 2015, the year neighborhood policing began (NYCLU 2019). In the fall of 2014, the Black Lives Matter movement launched protests in New York to demand racial justice and police reform. The protests were sparked by police killings of civilians across the country and by the refusal of a local grand jury to indict the NYPD officer who killed Eric Garner (Moynihan 2014). That December, many NYPD officers responded to the protests and

to a killing of two NYPD officers by slowing their work and reducing low-level arrests (Sullivan & O’Keeffe 2017). Though these high-profile events largely occurred before neighborhood policing was adopted, they require we isolate the effects of neighborhood policing from larger shifts. We exploit neighborhood policing’s staggered adoption to do this, as described below.

Through our analyses, we find that neighborhood policing had its strongest effect on low-level arrests, but not in the direction we expected. The typical precinct made 21% fewer proactive arrests and 12% fewer misdemeanor arrests the month neighborhood policing was implemented than our models estimate would have happened absent the policy. We observe no effect of neighborhood policing on violent crime, property crime, or racial disparities in arrest. The policy increased complaints of police misconduct, though only temporarily.

Neighborhood policing is changing law enforcement in two of the United States’ largest cities and is poised to spread elsewhere. If policymakers are to understand the impacts of this new strategy as they consider adopting it, it needs rigorous assessment. Here, we employ high-dimensional time series methods to reveal the policy’s impacts and guide lawmakers.

Extant research

Though no research has yet examined neighborhood policing specifically, scholarship on community policing provides some guidance as to how effective neighborhood policing’s community engagement strategies might be. Past studies of community policing have focused on two outcomes: crime and community attitudes toward the police. Research into community policing’s impact on arrests and racial disparities is less common.

Community engagement’s early proponents hoped the practice would decrease crime, improve residents’ relationships with police, and improve the quality of life for community members (Goldstein, 1987; Cordner 1997). New York’s neighborhood policing policy zeroed in

on crime control, making it the “first goal” of the program, followed by “promoting trust” and “solving problems collaboratively” (NYPD, 2018, 4). Two theories link community engagement and crime reduction. The first relies on improving police-community connections, and that will, in turn, enhance informal social control. By meeting with neighborhood groups and assisting them with problem solving, police empower “parochial” networks (e.g. neighbors, churches, and non-profit organizations) to surveil their areas, discourage crime, and regulate behavior (Reisig, 2010, 36). The second, related, mechanism reduces crime by increasing the amount of information residents relay to police (Braga, Brunson, & Drakulich, 2019). As New York Police Commissioner Bratton put it: “[W]here pockets of violence and crime remain in New York City, the NYPD will rely heavily on what these [new neighborhood coordination] officers can learn from their new community contacts about specific crimes, broader crime patterns, and general criminal activity” (Bratton, 2015, 4). Enhanced police-community communication and improved police legitimacy will, it is hoped, engender greater witness cooperation and information sharing.

Despite these theories, community engagement has had limited success in decreasing crime (Committee on Proactive Policing 2018). In a meta-analysis of 25 evaluations that included rigorous pre- and post-test designs, Gill et al. (2014) found no link between community engagement and serious crime. Studies that measured violent crime and property crime separately or analyzed many cities at once have similarly revealed no crime reduction benefits (Segrave & Collins 2005; MacDonald 2002). In an analysis of whether community policing improved the rate at which police agencies solve crimes, Roberts & Roberts (2016) found no effect of the practice on case clearance rates. Several studies have found community engagement reduced crime, however (Connell, Miggans, and McGloin 2008). A study with a randomized design found that when officers in Brooklyn Park, Minnesota spent their excess time building relationships with

community members at crime hot spots, crime in those areas decreased (Weisburd et al. 2020). Though the intervention in Minnesota was different from New York City's neighborhood policing (it included no specialized officers, it was concentrated at hot spots, the police force size was several hundred times smaller), the intervention shows community engagement can affect crime.

In addition to its crime control goal, New York City's neighborhood policing initiative aims "to promote trust and respect" between police and civilians (O'Neill, 2018, 4). Early in community policing's history, researchers worried that by putting police in increased contact with residents, the program would exacerbate historical tensions, increase police discrimination, and lead to more complaints of police misconduct from civilians (Moore 1992). Studies in Philadelphia and Houston, however, found community policing officers received either the same number or fewer civilian complaints of misconduct than traditional patrol officers (Hickman, Piquero, & Greene 2000; Kessler 1999). Repeated surveys of Chicago residents also indicated that the city's ambitious 1990s community policing initiative increased satisfaction with police (Skogan 2000). Similarly, a meta-analysis found people who live in areas that received community policing reported more confidence in police (Gill et al. 2014).

Past evaluations of community engagement have rarely considered the practice's impact on arrests, perhaps because community engagement programs rarely recommend changes in arrest practices (Chappell, MacDonald, & Manz 2006). Yet, a policy that alters how police interact with community members is likely to change how officers make arrests, especially discretionary ones. Community engagement might encourage low-level arrests if the increased contact exposes officers to more opportunities for arrest (Gascón & Roussell 2019). Because officers in the NYPD are trained in a broken-windows style of enforcement that prioritizes misdemeanor arrests, adding community engagement might invigorate this underlying approach (Davis, Mateu-Gelabert, &

Miller 2005; Gilmore & Gilmore 2016). Also, the increased time officers spend walking the beat, and the decreased time they spend responding to 911 calls, might shift their arrest energy toward the low-level arrests that stem from police observation and away from the more serious crimes that stem from victim reports. Alternatively, arrests might decrease during community engagement efforts if officers are disinclined to arrest community members after getting to know them better. Furthermore, the adoption of neighborhood policing might signal that the department's upper management is not judging officer quality solely by arrest productivity.

Which of these outcomes is most likely is difficult to predict because research on the community policing-arrest relationship is limited and findings vary by sample. Arrests declined in Cincinnati following a court-ordered implementation of community policing there (Ridgeway et al. 2009). At the national scale, two separate, cross-sectional analyses of around 200 cities each found no association between community engagement and arrests (Chappell, MacDonald, & Manz 2006; Eitle & Monahan 2009). The present study will build on this literature to analyze the New York City context with a longitudinal approach.

New York City's Police Commissioner singled out communities of color as especially likely to benefit from neighborhood policing (Bratton 2015, 2), though there is mixed evidence that past community engagement policies enhanced racial equity. In Chicago, White, Black, and Hispanic residents improved their opinions of police over the time community policing was implemented (Skogan 2009). However, community policing in Houston and Newark decreased fear of crime more among White residents than Black ones (Pate et al. 1986). There is ample evidence that misdemeanor policing in New York City has been racially disproportionate (Fagan et al. 2010; Beck 2020; Harcourt 2004). Our study will analyze whether neighborhood policing affected arrest disparities.

Data

To evaluate neighborhood policing's effects, we gathered data from four sources. Data on policing comes from the NYPD's Arrest Database, their Historic Complaint Database, and their Misconduct Complaint Dataset, the latter of which was released via the New York Civil Liberties Union following a lawsuit. The Census Bureau's American Community Survey (ACS) provides demographic information. Our unit of analysis is the precinct-month and we have outcome data on each of New York City's 76 residential precincts (excluding park and transit precincts) between January 2006 and September 2019, generating 165 months of data and 12,540 precinct-months.

Outcome Variables

We examine neighborhood policing's impact on four outcomes: crime, low-level arrests, racial disparities in low-level arrests, and complaints against police. The data for all four outcomes were originally address-level. We generate precinct-level counts by placing incidents in precinct boundaries using the city's precinct map shapefile and a geocoding program (Picard, 2015).

We then constructed a measure of each precinct-month's *violent crimes* and *property crimes*. The *violent crimes* variable is a count of all murders, manslaughters, felony assaults, and robberies made known to the police. Out of concerns for victim privacy, the NYPD does not provide information on the location of rapes, so they are omitted. *Property crimes* is a count of incidents of arson, burglary, grand larceny, petit larceny, and motor vehicle theft.

To capture the policy's impact on arrests, we constructed two measures of the low-level arrests. We focus on low-level arrests because they require the most police discretion and are therefore most susceptible to policy shifts. *Misdemeanor arrests* is a count of arrests for 133 misdemeanor offenses, the most common of which were marijuana possession, misdemeanor assault, theft of services (public transit fare evasion), possession of a controlled substance,

possession of stolen property, and trespassing. *Proactive arrests* are defined by the NYPD as those where “a substantial portion” of crimes are identified through police activity rather than victim complaints. They are for trespassing, drug possession, weapon possession, intoxicated driving, and possession of stolen property (NYPD 2015a, fn 1, 5, & 8). As a robustness check, we also model *felony arrests*, those for murder, manslaughter, felony assault, grand larceny, and burglary.

While New York City’s stop and frisk strategy received extensive attention, the practice was sharply curtailed beginning in 2012, and the reliability of New York stop counts has been in doubt following a federal court ruling limiting the practice, so we do not model stops in this project, though we do consider what impact the city-wide stop decline might have had on our outcomes.

One purpose of community engagement is to improve racial equity in policing. We analyze how neighborhood policing affected racial disparities in proactive arrests by constructing two measures of racial disproportionality. The first is the *difference between Black and White proactive arrest rates*. This is the rate at which White people are arrested per White precinct resident subtracted from the rate at which Black people are arrested per Black resident. The second racial disparity measure is the *Black-to-White proactive arrest rate ratio*. This is the Black rate divided by the White rate. Ratio measures can provide misleading results, particularly if the decision of which quantity to make the numerator is arbitrary (Stolzenberg 2021). Difference measures are less susceptible to this problem of non-reciprocity, so we use the difference measure as our primary outcome. However, in time-series contexts, difference measures can provide a skewed trend line because they reflect both the disparity in arrests and temporal changes in the aggregate number of arrests. Ratio measures eliminate this aggregate time trend. For this reason, we present both the difference and ratio measures in our descriptive graphs. The counts of residents by race used in each measure come from the ACS. Past research indicates that among all ethno-racial groups,

policing practices of Black and White people diverge the most (Weitzer, 2014), so those are the groups we compare here. We also modeled disparities in misdemeanor arrests and disparities between White arrests and Black and Latino arrests, and the results were similar. Because we do not have data on how frequently members of each racial group have contact with police nor data on racially specific offending rates, our approach cannot indicate racial bias, only racial disparity (Ridgeway & MacDonald 2010; Neil & Winship, 2019). Observed disparities might be due to differential contact, differential offending, or police bias.

We wanted to test whether neighborhood policing impacted *complaints of police misconduct*, a measure past work has used as a proxy for police-community relations (Kessler 1999). The variable is a count of accusations of police misconduct filed with the city's Civilian Complaint Review Board. We include accusations that were both substantiated and not substantiated by the Review Board because we are looking to capture public sentiment, not police misconduct. The most common complaints were for excessive use of force, cursing at a civilian, and making an unreasonable stop.

We also collected data on calls to 311, New York City's non-emergency hotline. However, in the simulation studies described below, 311 call data generated inaccurate predictions, due in part to a lack of the data's stationarity. Because we cannot be sure estimates of 311 data would be accurate, we do not model them here. New York City does not release data on 911 calls, so we are also unable to analyze those.

Treatment

The NYPD publicly announced the dates each precinct finished implementing neighborhood policing, and we use this announcement to calibrate the policy's implementation period (NYPD 2018). In the absence of information on when implementation began, we assume it

lasted one month, ending with the publicly announced date. Sensitivity analyses lagging the implementation dates backward two and three months revealed similar results to the one-month lag, changing only the timing of effects. These sensitivity analyses also give us confidence officers did not anticipate implementation and start behaving differently prior to the policy's adoption. The first precincts finished adopting the policy in May 2015, and the final ones in late 2018. City officials did not explain why some precincts received the treatment earlier than others, but analyses revealed no evident pattern, and our approach does not rely on randomness of treatment initiation time, only on stationarity of the outcomes.

A limitation of this study is that we do not have granular treatment implementation data. We do not know, for instance, how many meetings officers had with community members or whether “steady-sector” officers were truly given a third of their shifts away from responding to calls for service. We do have blunt implementation data. Through press accounts, we know each precinct received an average of 17 new neighborhood coordination officers, each restructured their precincts into smaller sectors, and each held regular community meetings and “build the block” collaboration events (Goodman 2015; Southall 2020; Sisak 2020).

Heterogeneous Treatment Effects

A host of criminological research has indicated that crime and policing vary across different neighborhood types. To measure the varying impact of neighborhood policing across heterogeneous precincts, we gathered demographic data from the ACS's tract-level estimates 2010-2014, the last time period before neighborhood policing's initiation. We combined tract data into precincts by using QGIS mapping software and city-provided shapefiles to determine the percent of each tract that overlapped with each precinct, then allocated the tract values to precincts in the same proportion. We analyze how the treatment varies by three measures of precincts' socio-

economic characteristics: their *unemployment rate*, the *percent of residents living in poverty*, and the *percent of residents without a B.A. degree*. To capture ethno-racial heterogeneity, we measure each precincts' *percent Black* residents and *percent Latino* residents. We use the *percent men aged 15-34* to represent the age and gender demographics most targeted by police and most commonly perpetrators of crime. As a measure of disorder, we capture the *percent of vacant housing units*.

Analytic Strategy

We are interested in the difference between what would have happened had the city not initiated neighborhood policing and what did happen after they initiated it. The difference between these two quantities is the causal effect of interest, or estimand. We fully observe what happened after neighborhood policing, but what would have happened absent the policy is an unknown counterfactual quantity we must estimate.

To do this, we let the effect of neighborhood policing for precinct i at q time points after treatment initiation be defined as $\Delta_{i,q}$. We let $q=0$ correspond to the first month neighborhood policing was implemented, let $q=1$ correspond to the time point one month after initiation, etc. One policy-relevant estimand is the average of this value over all precincts for a given q value, defined by $\Delta(q) = \frac{1}{n} \sum_{i=1}^n \Delta_{i,q}$. This shows the average impact of the policy q time points after it is initiated. We can examine this quantity for various q values to see how the effect of the policy varies over time. There could be a lag before the policy becomes effective, or conversely there could be a diminishing effect over time of such a policy as police revert back to their previous tendencies.

This quantity represents a *marginal effect* of the policy because it measures the average impact over all precincts in New York City. Also of interest are heterogeneous effects that tell us whether the impact of the policy varies across precinct-level characteristics. To do this, we can calculate the correlation between $\Delta_{i,q}$ and observed characteristics in our data, denoted by

$\rho(X_{ij}, \Delta_{i,q})$. This will inform us of how the covariates relate to the impact of the policy. We can examine the magnitude of these correlations to understand how each covariate relates to the treatment effect, and we can look at their confidence intervals to determine if any covariates are significantly associated with the size of the treatment effect. For instance, if we find a positive correlation between the treatment effect and the percentage of residents in a precinct without a bachelor's degree, that means neighborhood policing had more positive impacts in precincts with lower numbers of college-educated people.

Having discussed the estimands of interest, we now discuss how we estimate them. All of the quantities described above are functions of $\Delta_{i,q}$. Once we have estimates of $\Delta_{i,q}$, it is straightforward to estimate the policy relevant estimands defined by $\Delta(q)$, and $\rho(X_{ij}, \Delta_{i,q})$. For the marginal estimands, we simply sum the relevant estimates of $\Delta_{i,q}$, while for the heterogeneous treatment effects we calculate the Pearson correlation coefficient between the estimated $\Delta_{i,q}$ values and each covariate. For heterogeneous treatment effects we focus on correlations between the covariates and the estimated treatment effects in the first three time periods post treatment initiation as estimates are the most reliable in time periods closest to the adoption of the policy. In light of this, the fundamental problem left is how to estimate $\Delta_{i,q}$.

We first describe the idea in the simpler setting of only one precinct. Methods for settings with one unit have been described in Broderson et al. (2015) and Papadogeorgou et al. (2018). This method fits a Bayesian time series model to the data from the period before initiation of neighborhood policing to forecast what would have happened had neighborhood policing never been initiated. An illustration of our approach using simulated data can be found in Figure 1. The black line shows the time series before treatment, at time 70. The red line shows the time series after the policy has been initiated, while the dotted line shows our model's forecast for what would

have happened in the absence of the policy. The gray region represents 95% credible intervals. The difference between the red line and the dotted line is our best estimate of the effect of the policy, and the uncertainty in this estimate is fully characterized by the size of the shaded region.

[Figure 1 at end of document]

This illustration demonstrates how to estimate $\Delta_{i,q}$ for a single precinct. However, we observe 76 time series, one for each precinct in New York City. A complicating factor is that the 112 to 150 time points pre-policy implementation for each precinct generates a high-dimensional data set with the number of units similar to the number of time points. To address this, we use a high-dimensional Bayesian vector autoregressive model that accounts for both temporal autocorrelation and spatial correlation across precincts at a given point in time. Letting Y_t be the vector of outcomes observed for all 76 precincts at time t , we posit the following model:

$$Y_t = \mu(t) + A(Y_{t-1} - \mu(t-1)) + \epsilon_t, \epsilon_t \sim MVN(0, \Sigma).$$

Here, $\mu(t)$ is a vector of mean functions evaluated at time t . This can be structured to account for trends across time or for seasonality. We use natural splines to model this function and the degrees of freedom are chosen by the validation approach described below. For each precinct, we estimate a separate mean function that allows for both a precinct-specific intercept and a precinct-specific trend across time, which will account for any unmeasured time-invariant covariates. A is a 76-by-76 matrix of coefficients that determines how the residuals at time $t-1$ affect time t . This serves two purposes: it can improve predictions by using data from the previous time point to guide forecasts, and it helps to account for temporal autocorrelation. Ignoring such temporal autocorrelation could lead to inflated type I error rates and overly small confidence intervals. Lastly, the errors ϵ_t allow for correlation across precincts at a given point in time. This correlation is driven by the covariance matrix Σ , which is also unknown and must be estimated.

This model cannot be fit without imposing constraints on the unknown parameters, as there are currently more parameters than samples in the data. It is likely that neighboring precincts are more correlated than distant precincts. With this in mind, we impose structure on the covariance matrix Σ by assuming that precincts are conditionally independent of each other if they are not neighbors, when conditioning on their neighboring precinct values. This can easily be enforced by restricting the (i,j) elements of Σ^{-1} to be zero if precincts i and j are not neighbors. This approach draws on the techniques in Antonelli & Beck (2021). A similar strategy will be used for the A matrix, by forcing elements (i,j) of A to be zero if precincts i and j are not neighbors.

These approaches to reducing the dimension of both A and Σ allow us to proceed with updating all remaining parameters using Markov chain Monte Carlo sampling. We use an empirical Bayes approach to estimating Σ under the constraints above, and non-informative normal prior distributions for all remaining components of A and $\mu(t)$. Once we have the posterior distribution of all unknown parameters in our model, we can easily produce forecasts of future time points along with corresponding measures of uncertainty. Obtaining the posterior distribution of $\Delta_{i,q}$ is straightforward once we have the posterior distribution of forecasts for what would have happened in the absence of the policy. For every posterior draw, we simply take the observed data post-treatment initiation and subtract the posterior draw of the forecast from our model, and this difference is a posterior draw of $\Delta_{i,q}$.

While the aforementioned model can be used to produce estimates and confidence intervals for the estimands of interest, our results are contingent on correct specification of this model, as well as stationarity of the outcome time series. If our model does not fit the data well, we could obtain biased estimates or we could have incorrect confidence interval widths that lead to inflated type I error rates. Fortunately, in time series settings we can empirically evaluate whether the

assumptions required for our approach are likely to hold. We pretend we observe fewer time points than we actually do, then use our model to predict the future time series and compare our predictions to the ground truth, the data we pretended not to observe. We do this many times, each time varying the number of time points we observe and tracking our model's performance. This form of validation was used in Schell et al. (2018) to evaluate the effectiveness of gun control legislation. We ran the simulation for each outcome, and our approach provided relatively unbiased estimates of treatment effects while also generating accurate estimates of uncertainty such that the 95% confidence intervals cover the true values in 95% of the simulated data. A simulation of 311 call data revealed they do not produce accurate estimates in this framework, likely due to the data's lack of stationarity, so we omit this outcome from our results.

Results

Descriptive results

Figure 2 displays the city-wide trends over time for each of our primary outcome variables, measured monthly. The graphs cover different time periods depending on data availability. Neighborhood policing was implemented between May 2015 and October 2018. The top graph shows aggregate property and violent crime fluctuated seasonally and declined during the study window. Our simulations indicate that the within-year, seasonal variation noticeable on this graph is accounted for by our model. Misdemeanor arrests also declined during these years, with the subset of misdemeanor arrests that were proactive declining even more sharply. The difference between the number of arrests of Black people per Black resident and the number of arrests of White people per White resident declined during the study window. As mentioned above, difference measures include changes both in the racial disparity and in the aggregate arrest count, which can produce misleading trends in a descriptive context. Here, much of the decline in the

difference measure is due to the decline in arrests for all groups, not a change in the racial disparity. The next graph corrects for this by portraying the ratio measure, which eliminates the aggregate trend. This graph indicates that, in January 2010, the NYPD's arrest rate of Black people was 5.7 times that of White people. The disparity declined steadily, then leveled off in 2017 such that the ratio in December 2018 was 3.4. Complaints of police misconduct were volatile, and they rose somewhat after 2017. While no causal conclusions can be drawn from these descriptive, city-wide graphs, they paint a picture of a city experiencing a gradual decline in crime, a sharp decline in low-level arrests, and stasis in both the racial proportionality of proactive arrests and in complaints against police.

[Figure 2 at end of document]

Modeling Results

Figure 3 displays the marginal effects of neighborhood policing for each outcome across all precincts. The vertical line on each graph represents the implementation month of the policy, with each precinct's time series centered around this zero date. The solid lines are the observed data, both before and after implementation. They differ from the descriptive graphs presented in Figure 2 because they reflect average values between precincts rather than total city values and are recentered around each precinct's unique implementation month. The dotted lines in Figure 3 represent what the models estimate would have happened absent neighborhood policing. The difference between what was observed and estimated after implementation represents the effect of neighborhood policing. The grey bands are 95% credible intervals around the estimates, so any observed value outside those bands is a statistically significant effect of the policy.

Neighborhood policing did not influence either violent or property crime at a statistically significant level at any time point we observed. Because the estimated values' confidence bands

contain the observed values at each time point, we cannot reject the null hypothesis that neighborhood policing does not affect these outcomes.

Complaints of police misconduct experienced a statistically significant increase in the four to seven months after neighborhood policing's implementation. The typical precinct experienced about 2.3 complaints in the fifth month, 35% more than the expected 1.9 complaints. However, in month eight, observed complaints recross the confidence interval, so the effect appears to be temporary. This effect is largely driven by an increase in complaints in the fifth and sixth months after implementation. Such a brief effect provides only weak evidence that neighborhood policing impacted complaints against police. We performed a post hoc analysis to determine what kind of complaints drove this effect. These results are displayed in Appendix B and indicate it was complaints of "discourtesy/abuse of authority" such as allegations of officers cursing at residents or making unwarranted stops that drove this temporary increase, not complaints of excessive force.

Neighborhood policing had a large effect on arrests. The plots of proactive and misdemeanor arrests each show a sharp difference in the observed trendlines before and after implementation, and these observed arrests are statistically significantly lower than what we estimated would have happened absent neighborhood policing. The typical precinct made 46 proactive arrests the month it implemented neighborhood policing, 21% fewer than the 58 arrests our model predicted and 16% fewer than the 55 made the month before. That immediate reduction maintained through each subsequent month, with the typical precinct making 37% fewer proactive arrests than they otherwise would have in the ninth month after adoption. Misdemeanor arrests experienced a similarly sized decline, though it did not persist for as long. The average precinct made 165 misdemeanor arrests the month of implementation, 12% fewer than the expected 187 and 9% fewer than the 182 made the month before. Five months out, arrests were still lower than

expected, but not at a statistically significant level. This reduction, though not statistically significant after month four, endured through the ten months measured here.

Looking just at the solid lines of observed arrest data and ignoring for a moment the dotted lines estimating what would have happened absent neighborhood policing, the difference in both arrest types the month before implementation and the month of implementation is stark. A sharp decrease in arrests occurred the month neighborhood policing was introduced to each precinct, and since each precinct adopted the policy at different times, this break supports the conclusion from the modeling that it is neighborhood policing—and not a city-wide change like protests, an NYPD work slowdown, or the continued decline in stop and frisk—that drove the arrest decline. To see if this decline was a uniquely low-level arrest phenomenon, we analyzed felony arrests and found that though these less discretionary arrests declined the month of implementation, they were higher than expected by month four. Those results are presented in Appendix B.

While neighborhood policing reduced the number of low-level arrests officers made, it did not change the racial proportion of such arrests. The bottom-right plot in Figure 3 shows the impact of neighborhood policing on the difference between the Black and white proactive arrest rates. This outcome is notably consistent over time, and the values observed following implementation were either identical to the estimated values or well within the estimate's margins of error.

[Figure 3 at end of document]

As a sensitivity analysis, we compared our results with those from interrupted time series models. Our primary approach relies on the assumption that the outcome's time series is stationary, while the interrupted time series approach relies on the assumption that there are no unmeasured time-varying confounders. These two approaches rely on different model specifications as well as different assumptions, so similar results across each approach would suggest our findings are not

artifacts of the chosen statistical approach, but true effects of neighborhood policing. Indeed, each approach found very similar results. The results for crime, proactive arrests, and racial disparities were substantively identical. The timing of the effects for civilian complaints against police and misdemeanor arrests changed by one or two months but were otherwise the same. Full results are presented in Appendix A. This sensitivity analysis gives us further trust in our findings.

The impact of neighborhood policing might not have been uniform across different precincts. To account for possible heterogeneous treatment effects, we estimate how the policy's influence varied across seven precinct characteristics in the first three months after implementation. Those results are represented by the heat grid in Figure 4. A negative effect of neighborhood policing in some types of precincts and a positive effect in others could have cancelled out to obscure a true effect of the policy. We were especially interested in whether this happened with the outcomes where we observed no aggregate effect: crimes and racial disparities. Figure 4 reveals that the effect of neighborhood policing on those two outcomes and on complaints against police did not vary by precinct characteristics. The effect on arrests, however, did. Neighborhood policing had a stronger suppressive effect on misdemeanor arrests in precincts home to more Black residents and it had a stronger suppressive effect on proactive arrests in precincts with more Latino residents. In other words, neighborhood policing reduced arrest rates more in those precincts than it did in the typical precinct. The effect of neighborhood policing on arrests also varied by the three socio-economic variables. Precincts with higher unemployment, poverty, and non-college-education rates experienced a stronger negative effect of neighborhood policing on both types of arrests. The policy reduced arrests more in poorer neighbors and reduced arrests less in wealthier ones.

[Figure 4 at end of document]

Another question of interest is where in New York City neighborhood policing had the strongest effects. Because neighborhood policing had the largest impact on proactive arrests, we estimated the treatment effect on proactive arrests for each precinct in the three months after implementation. Figure 5 maps these precinct-specific estimates, with darker shades representing stronger negative effects of neighborhood policing on proactive arrests. The darkest shade indicates precincts where our models estimate police made between 50 and 75 fewer proactive arrests than they otherwise would have in the three months after neighborhood policing was initiated. The lighter shades mark precincts where neighborhood policing was estimated to increase the number of arrests.

The map shows the policy reduced arrests the most in upper Manhattan, the Bronx, and the precincts of Brooklyn and Queens farthest from Manhattan. A comparable map showing the effects of misdemeanor arrests, omitted here for space, showed similar spatial patterns. With the exception of the 7th precinct on Manhattan’s Lower East Side, the precincts that saw the largest reductions in arrests because of neighborhood policing were outside of the lower- and midtown-Manhattan areas that host the City’s central business district, cultural centers, and tourist destinations. Reading this map along with the heat grid in Figure 4 reveals that neighborhood policing’s suppressive effect on low-level arrests was strongest in the higher-unemployment, higher-poverty, less-college-educated outer boroughs.

[Figure 5 about here]

Discussion and Conclusion

In 2015, New York City announced “neighborhood policing,” an initiative aimed at reducing crime and improving police-community relations. The policy restructured the city’s precincts, appointed hundreds of new “neighborhood coordination” officers, and established

“steady-sector” officers who walked a consistent beat and had a third of their shifts away from emergency response to address community problems. We use high-dimensional time series models to estimate neighborhood policing’s effects. Our approach accounts for the policy’s staggered adoption, it addresses potential correlation among outcomes and across precincts, and it examines treatment effect heterogeneity across precincts. We also validated our method using simulations with observed data from before the initiation of the policy.

New York City’s leadership promoted neighborhood policing first as a crime control strategy. We found the program did not reduce property or violent crime, at least in the first ten months after implementation. This aligns with past research into other community engagement programs. We found the policy increased complaints of police misconduct, but only temporarily. The racial equity aspirations of neighborhood policing were not realized, as the racial disparity in low-level arrests was unchanged.

We had expected that the increased contact between police and community members might accentuate the longstanding misdemeanor-focused practices of the NYPD and lead to more low-level arrests. The opposite proved true. The policy reduced misdemeanor arrests for at least four months after its adoption and proactive arrests for the entire ten months analyzed. The arrest reduction finding was particularly striking because of the increase in the number of police officers during the study period. Adding more potential arresters did not result in more arrests. Analyses of heterogeneous treatment effects revealed the policy’s impact on low-level arrests was especially strong in poorer neighborhoods outside Manhattan.

While our methods cannot reveal the mechanisms linking neighborhood policing with decreased arrests, there are several possibilities. We do not suspect the policy decreased arrests by decreasing the underlying low-level crime rate because the policy did not affect violent or property

crime rates. Assuming misdemeanor crime rates and serious crime rates move in tandem, neighborhood policing left misdemeanor crime unchanged, and the arrest effect was driven by police choosing to make fewer arrests rather than there being fewer behaviors for which to make an arrest. This hypothesis is further supported by our finding that the policy had no consistent impact on felony arrests, which are less discretionary.

Perhaps officers made fewer low-level arrests because they were dissatisfied with the new neighborhood policing policy, and they initiated a work slowdown to protest it. This is something NYPD officers have done when they disagree with past city government decisions (Sullivan & O’Keeffe 2017). This is an unlikely explanation for several reasons. First, our models observed a decline in arrests localized around the time-staggered implementation of neighborhood policing, but the NYPD slowdown affected all precincts at the same time. Second, journalistic accounts found rank-and-file officers were largely supportive of, or at least neutral toward, neighborhood policing (Goodman 2015). Finally, the president of New York’s patrol officers’ union was not critical of neighborhood policing as he was of local government decisions during previous slowdowns (Lynch 2016).

Black Lives Matter protests and the NYPD’s declining use of stop-and-frisk were other high-profile events during our study window. As with the work slowdown, they affected all precincts simultaneously, but the effects we observe were following the month neighborhood policing was implemented, meaning it was not these city-wide changes driving these findings.

One possible explanation for the decrease in arrests is that officers changed their attitudes toward arrest. The NYPD management’s adoption of neighborhood policing might have signaled to street-level officers the department was deprioritizing arrests as a performance metric. One press account found officers who had previously focused on arrests were the same ones making

community visits under neighborhood policing (Goodman, 2015). Relatedly, the policy might have made officers less inclined to arrest community members for whom they now had a greater understanding. This is speculation, as there is nothing in our data to suggest this mechanism, but future interviews with officers implementing neighborhood policing could tease out the connection between the policy and the arrest decline. While changing arrests is rarely a stated goal of community engagement policies, these results suggest it can be a prominent effect.

The increase in complaints against police that we observe was very temporary, so it ought not be overinterpreted, but we can think of two reasons why neighborhood policing might be having such an effect. First, the increased contact with residents that neighborhood policing engenders might increase opportunities for negative interactions. Even if the proportion of police-civilian interactions that generate complaints did not change, their growing total number might increase complaints. Second, neighborhood policing grants officers autonomy to solve problems. Both the increased contact and discretion might exacerbate underlying problems like use of excessive force and discriminatory policing where those are prevalent, especially in a city like New York that historically prioritized intense enforcement.

The present study has some limitations. First, as mentioned above, we do not have granular implementation data. We think neighborhood policing is an important enough policy to evaluate even in the absence of precise information on how many meetings were held or how faithfully precincts gave officers time away from 911 call response. Also, we have blunt implementation data, including the dates each precinct received 17 new officers and reorganized into smaller sectors. If the city weakly adopted neighborhood policing, this study's findings reflect the impacts of the changes that were made, like hiring new officers, more than the changes that were only

promised. We hope future research, and the city's own long-promised internal evaluation of neighborhood policing (Smith 2020), will have more detailed implementation data.

A second limitation is our inability to control for crime reporting inflation (Weisburd et al. 2020). If neighborhood policing was effective at both reducing crime and encouraging the public to report more crimes, these two changes might neutralize one another and make it appear as though crime remained the same when it declined. This is a problem with almost all research on community engagement policies. Conducting victimization surveys in each precinct before and after the policy was implemented would capture the true crime rate, but such an approach is expensive, especially in a city as large as New York. To address potential crime reporting inflation within our means, we analyzed property crime, which is strongly susceptible to reporting bias, and violent crime, which is much less so. For both crime outcomes, we found the same results, providing some reassurance that increased reporting is not biasing our results.

Finally, we are limited by our time horizon. It is possible neighborhood policing will take longer than ten months, the extent of our analysis, to have an impact on crime and other outcomes. This is mostly unavoidable. Analyzing the impacts of a policy years out from its adoption poses serious methodological challenges as it becomes difficult to separate effects of the policy from other intervening changes. Nevertheless, it is important to acknowledge that our study only applies to a short time window after adoption.

To our knowledge, this is the first evaluation of New York City's neighborhood policing initiative. Despite a lack of evidence on the initiative's efficacy, Chicago launched the program in the summer of 2020, modeling their approach "off the Neighborhood Policing Structure in New York City" (Office of the Mayor of Chicago, 2020). As policymakers continue to pursue neighborhood policing, our results recommend caution. We find the policy is more effective at

changing police behaviors like arrest frequency than reducing community characteristics like crime rates or complaints against police.

Neighborhood policing in New York did have a strong suppressive effect on low-level arrests, and the arrest declines happened without increasing crime. Elected officials might want to pursue neighborhood policing as a means of reducing low-level arrests. Doing so is a laudable goal since such arrests have stark, negative consequences for the people arrested and for police-community relations (Natapoff 2018; Fagan, Tyler, & Meares 2016; Justice 2021). Additionally, we found neighborhood policing's arrest reduction was greatest in poorer neighborhoods, areas traditionally exposed to the highest levels of discretionary arrests. Yet, policymakers should be cautious about adopting this policy as an arrest-reduction strategy because there are likely other, less expensive, ways to reduce low-level arrests, and we did not find that the racial disparity in arrests changed. Furthermore, if a city increases its force capacity with neighborhood policing, the new officers could be redirected to more arrest-intensive strategies by future mayors, the way New York City's community policing policy of the 1980s set the stage for the City's broken-windows arrest increases in the 1990s (Beck & Matles 2015).

We find New York City's neighborhood policing initiative did not reduce crime rates, complaints against police, or racial equity in low-level arrests. It did reduce aggregate low-level arrests in a persistent way, a not insignificant feat. As policymakers consider adopting neighborhood policing, they might reflect on New York's case to understand the policy's limits.

References

- Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(490), 493–505. <https://doi.org/10.1198/jasa.2009.ap08746>
- Antonelli, J., & Beck, B. (2021). Estimating heterogeneous causal effects in time series settings with staggered adoption: An application to neighborhood policing. *ArXiv:2006.07681 [Stat]*. <http://arxiv.org/abs/2006.07681>
- Ashenfelter, O. (1978). Estimating the Effect of Training Programs on Earnings. *The Review of Economics and Statistics*, 60(1), 47. <https://doi.org/10.2307/1924332>
- Beck, B. (2020). Policing Gentrification: Stops and Low-Level Arrests during Demographic Change and Real Estate Reinvestment. *City & Community*, 19(1), 245–272. <https://doi.org/10.1111/cico.12473>
- Beck, B., & Matles, A. (2015, June 24). We need fewer NYPD officers — not more. *AM New York*. <https://www.amny.com/opinion/we-need-fewer-nypd-officers-not-more-1-10575625/>
- Braga, A. A., Brunson, R. K., & Drakulich, K. M. (2019). Race, Place, and Effective Policing. *Annual Review of Sociology*, 45(1), 535–555. <https://doi.org/10.1146/annurev-soc-073018-022541>
- Bratton, W. J. (2015). *The NYPD Plan of Action and the Neighborhood Policing Plan: A Realistic Framework for Connecting Police and Communities*. New York Police Department. <http://www.nyc.gov/html/nypd/html/home/POA/pdf/Plan-of-Action.pdf>
- Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N., & Scott, S. L. (2015). Inferring causal impact using Bayesian structural time-series models. *The Annals of Applied Statistics*, 9(1), 247–274. <https://doi.org/10.1214/14-AOAS788>
- Chappell, A. T., MacDonald, J. M., & Manz, P. W. (2006). The Organizational Determinants of Police Arrest Decisions. *Crime & Delinquency*, 52(2), 287–306. <https://doi.org/10.1177/0011128705278329>
- Committee on Proactive Policing: Effects on Crime, Communities, and Civil Liberties, Committee on Law and Justice, Division of Behavioral and Social Sciences and Education, & National Academies of Sciences, Engineering, and Medicine. (2018). *Proactive Policing: Effects on Crime and Communities* (David Weisburd & M. K. Majimundar, Eds.). National Academies Press. <https://doi.org/10.17226/24928>

- Connell, N. M., Miggans, K., & McGloin, J. M. (2008). Can a Community Policing Initiative Reduce Serious Crime?: A Local Evaluation. *Police Quarterly*, 11(2), 127–150. <https://doi.org/10.1177/1098611107306276>
- Cordner, G. W. (1997). Community Policing: Elements and Effects. In R. G. Dunham & G. P. Alpert (Eds.), *Critical Issues in Policing* (pp. 451–468). Waveland Press.
- Davis, R. C., Mateu-Gelabert, P., & Miller, J. (2005). Can Effective Policing Also Be Respectful? Two Examples in the South Bronx. *Police Quarterly*, 8(2), 229–247. <https://doi.org/10.1177/1098611104269531>
- Edelman, S. (2020, August 15). NYPD axes \$4.1M ‘Sentiment Meter’ cop-favorability poll. *The New York Post*. <https://nypost.com/2020/08/15/nypd-axes-4-1m-sentiment-meter-cop-favorability-poll/>
- Eitle, D., & Monahan, S. (2009). Revisiting the Racial Threat Thesis: The Role of Police Organizational Characteristics in Predicting Race-Specific Drug Arrest Rates. *Justice Quarterly*, 26(3), 528–561. <https://doi.org/10.1080/07418820802427817>
- Fagan, J. A., Geller, A., Davies, G., & West, V. (2010). Street Stops and Broken Windows Revisited. In S. K. Rice & M. D. White (Eds.), *Race, Ethnicity, and Policing: New and Essential Readings* (pp. 309–348). NYU Press. <https://doi.org/10.18574/nyu/9780814776155.003.0013>
- Fagan, J., Tyler, T., & Meares, T. L. (2016). Street Stops and Police Legitimacy in New York. In *Comparing the Democratic Governance of Police Intelligence*. Edward Elgar Publishing.
- Gascón, L. D., & Roussell, A. (2019). *The limits of community policing: civilian power and police accountability in black and brown Los Angeles*. New York University Press.
- Gill, C., Weisburd, D., Telep, C. W., Vitter, Z., & Bennett, T. (2014). Community-oriented policing to reduce crime, disorder and fear and increase satisfaction and legitimacy among citizens: a systematic review. *Journal of Experimental Criminology*, 10(4), 399–428. <https://doi.org/10.1007/s11292-014-9210-y>
- Gilmore, R. W., & Gilmore, C. (2016). Beyond Bratton. In *Policing the Planet: Why the Policing Crisis Led to Black Lives Matter*. Verso Books.
- Goldstein, H. (1987). Toward Community-Oriented Policing: Potential, Basic Requirements, and Threshold Questions. *Crime & Delinquency*, 33(1), 6–30. <https://doi.org/10.1177/0011128787033001002>
- Goodman, J. D. (2015, August 23). In New York, Testing Grounds for Community Policing. *The New York Times*. <https://www.nytimes.com/2015/08/24/nyregion/for-new-york-police-a-radical-change-for-queens-residents-a-step.html>
- Harcourt, B. E. (2004). *Illusion of order: the false promise of broken windows policing*.
- Hickman, M. J., Piquero, A. R., & Greene, J. R. (2000). Does Community Policing Generate Greater Numbers and Different Types of Citizen Complaints than Traditional Policing? *Police Quarterly*, 3(1), 70–84. <https://doi.org/10.1177/1098611100003001003>

- Kessler, D. A. (1999). The Effects of Community Policing on Complaints Against Officers. *Journal of Quantitative Criminology*, 15(3), 333–372. <https://doi.org/10.1023/A:1007580632179>
- Lopez Bernal, J., Cummins, S., & Gasparrini, A. (2016). Interrupted time series regression for the evaluation of public health interventions: a tutorial. *International Journal of Epidemiology*, dyw098. <https://doi.org/10.1093/ije/dyw098>
- Lum, C., & Nagin, D. S. (2017). Reinventing American Policing. *Crime and Justice*, 46(1), 339–393. <https://doi.org/10.1086/688462>
- Lynch, P. J. (2016, October 27). For Real Community Policing, Let Officers Do Their Jobs. *The New York Times*. <https://www.nytimes.com/2016/10/27/opinion/for-real-community-policing-let-officers-do-their-jobs.html>
- MacDonald, J. M. (2002). The Effectiveness of Community Policing in Reducing Urban Violence. *Crime & Delinquency*, 48(4), 592–618. <https://doi.org/10.1177/001112802237131>
- Moore, M. H. (1992). Problem-Solving and Community Policing. *Crime and Justice*, 15, 99–158. <https://doi.org/10.1086/449194>
- Moynihan, C. (2014, December 9). Defying Ban, Students March to Brooklyn in Protest of Eric Garner Decision. *New York Times*.
- Natapoff, A. (2018). *Punishment without crime: how our massive misdemeanor system traps the innocent and makes America more unequal* (First edition). Basic Books.
- Neil, R., & Winship, C. (2019). Methodological Challenges and Opportunities in Testing for Racial Discrimination in Policing. *Annual Review of Criminology*, 2(1), 73–98. <https://doi.org/10.1146/annurev-criminol-011518-024731>
- NYPD. (2015a). *Misdemeanor Offenses by Precinct.* Historical New York City Crime Data. Historical New York City Crime Data. <http://www1.nyc.gov/site/nypd/stats/crime-statistics/historical.page>
- NYPD. (2015b, June 26). *Mayor de Blasio, Commissioner Bratton Unveil New, Groundbreaking Neighborhood Policing Vision.* <https://www1.nyc.gov/office-of-the-mayor/news/444-15/transcript-mayor-de-blasio-commissioner-bratton-new-groundbreaking-neighborhood-policing>
- NYPD. (2018). *Shared Responsibility Viewbook: A New Yorker's Guide to Neighborhood Policing.* New York City Police Department. <https://cloud.3dissue.com/167796/168091/212990/Shared-Responsibility-Viewbook/index.html>
- Office of the Mayor of Chicago. (2020). *Mayor Lightfoot and CPD Expand Neighborhood Policing Initiative* [Press Release]. City of Chicago.
- Oliver, W. M. (2000). The Third Generation of Community Policing: Moving Through Innovation, Diffusion, and Institutionalization. *Police Quarterly*, 3(4), 367–388. <https://doi.org/10.1177/109861110000300402>
- O'Neill, J. P. (2018). *The Police Commissioner's Report 2018.* New York City Police Department.

- Papadogeorgou, G., Mealli, F., Zigler, C. M., Dominici, F., Wasfy, J. H., & Choirat, C. (2018). Causal Impact of the Hospital Readmissions Reduction Program on Hospital Readmissions and Mortality. *ArXiv:1809.09590 [Stat]*. <http://arxiv.org/abs/1809.09590>
- Pate, A. M., Wycoff, M. A., Skogan, W. G., & Sherman, L. W. (1986). *Reducing Fear of Crime in Houston and Newark: A Summary Report*. Police Foundation. <https://www.policefoundation.org/wp-content/uploads/2015/07/Pate-et-al.-1986-Reducing-Fear-of-Crime-in-Houston-and-Newark-Summary-Report-.pdf>
- Picard, R. (2015). *Geoinpoly: Stata module to match geographic locations to shapefile polygons*. <https://ideas.repec.org/c/boc/bocode/s458016.html>
- Reed, W. E. (1999). *The politics of community policing: the case of Seattle*. Garland Pub.
- Reisig, M. D. (2010). Community and Problem-Oriented Policing. *Crime and Justice*, 39(1), 1–53. <https://doi.org/10.1086/652384>
- Ridgeway, G., & MacDonald, J. (2010). *Methods for Assessing Racially Biased Policing*. NYU Press. <https://doi.org/10.7249/RP1427>
- Ridgeway, G., Schell, T., Gifford, B., Saunders, J., Turner, S., Riley, K., & Dixon, T. (2009). *Police-Community Relations in Cincinnati*. RAND Corporation. <https://doi.org/10.7249/MG853>
- Roberts, A., & Roberts, J. M. (2016). Crime Clearance and Temporal Variation in Police Investigative Workload: Evidence from National Incident-Based Reporting System (NIBRS) Data. *Journal of Quantitative Criminology*, 32(4), 651–674. <https://doi.org/10.1007/s10940-015-9270-9>
- Schell, T., Griffin, B. A., & Morral, A. (2018). *Evaluating Methods to Estimate the Effect of State Laws on Firearm Deaths: A Simulation Study*. RAND Corporation. <https://doi.org/10.7249/RR2685>
- Segrave, M. & Collins, L. (2005). *Evaluation of a suburban crime prevention team*. Australian Institute of Criminology. <http://www.aic.gov.au/publications/tbp/tbp014/tbp014.pdf>
- Sisak, M. (2020). *NYPD commissioner to bring new focus on community policing*. Associated Press. <https://www.police1.com/community-policing/articles/nypd-commissioner-to-bring-new-focus-on-community-policing-q9jQr21XADNKcWdl/>
- Skogan, W. G. (2009). *Police and Community in Chicago: A Tale of Three Cities*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195154580.001.0001>
- Skogan, W. G., & Hartnett, S. M. (2000). *Community policing, Chicago style*. Oxford Univ. Press.
- Smith, G. B. (2020, July 23). De Blasio’s ‘Stalled’ Pet NYPD Neighborhood Policing Plan Draws Outside Look. *The City*. <https://www.thecity.nyc/2020/7/23/21336402/de-blasios-nypd-neighborhood-policing-plan-draws-outside-look>
- Southall, A. (2020, August 24). Shootings Have Soared. Is the N.Y.P.D. Pulling Back? *The New York Times*. <https://www.nytimes.com/2020/07/16/nyregion/nyc-shootings-nypd.html>

- Sullivan, C. M., & O’Keeffe, Z. P. (2017). Evidence that curtailing proactive policing can reduce major crime. *Nature Human Behaviour*, *1*(10), 730–737. <https://doi.org/10.1038/s41562-017-0211-5>
- Stolzenberg, R. M. (2021). New results about difference scores, ratios, and hierarchical linear model parameters as tools for comparing groups. *Sociological Methods & Research*, *50*(2), 467–490. <https://doi.org/10.1177/0049124118789715>
- Weisburd, David, & Eck, J. E. (2004). What Can Police Do to Reduce Crime, Disorder, and Fear? *The ANNALS of the American Academy of Political and Social Science*, *593*(1), 42–65. <https://doi.org/10.1177/0002716203262548>
- Weisburd, David, Gill, C., Wooditch, A., Barritt, W., & Murphy, J. (2020). Building collective action at crime hot spots: Findings from a randomized field experiment. *Journal of Experimental Criminology*. <https://doi.org/10.1007/s11292-019-09401-1>
- Weitzer, R. (2014). *Police Race Relations*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199843886.013.020>

Figures

Figure 1. Toy example of approach

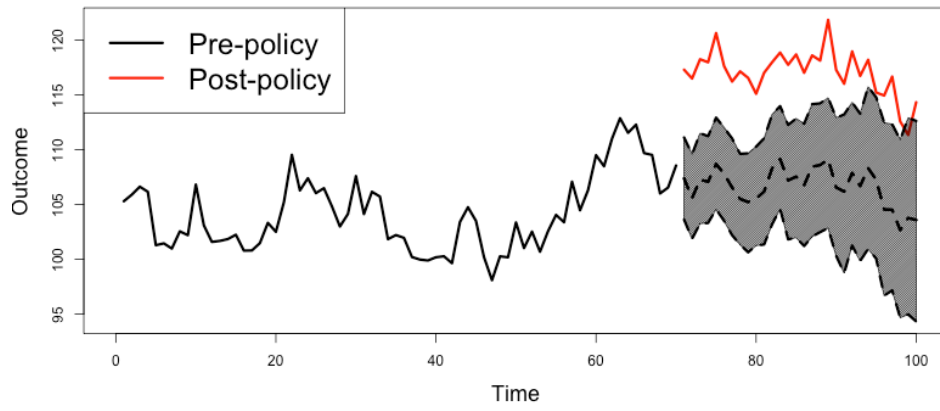


Figure 2. Outcome variables over time




Figure 3. Observed values and estimated values absent neighborhood policing, months before and after implementation

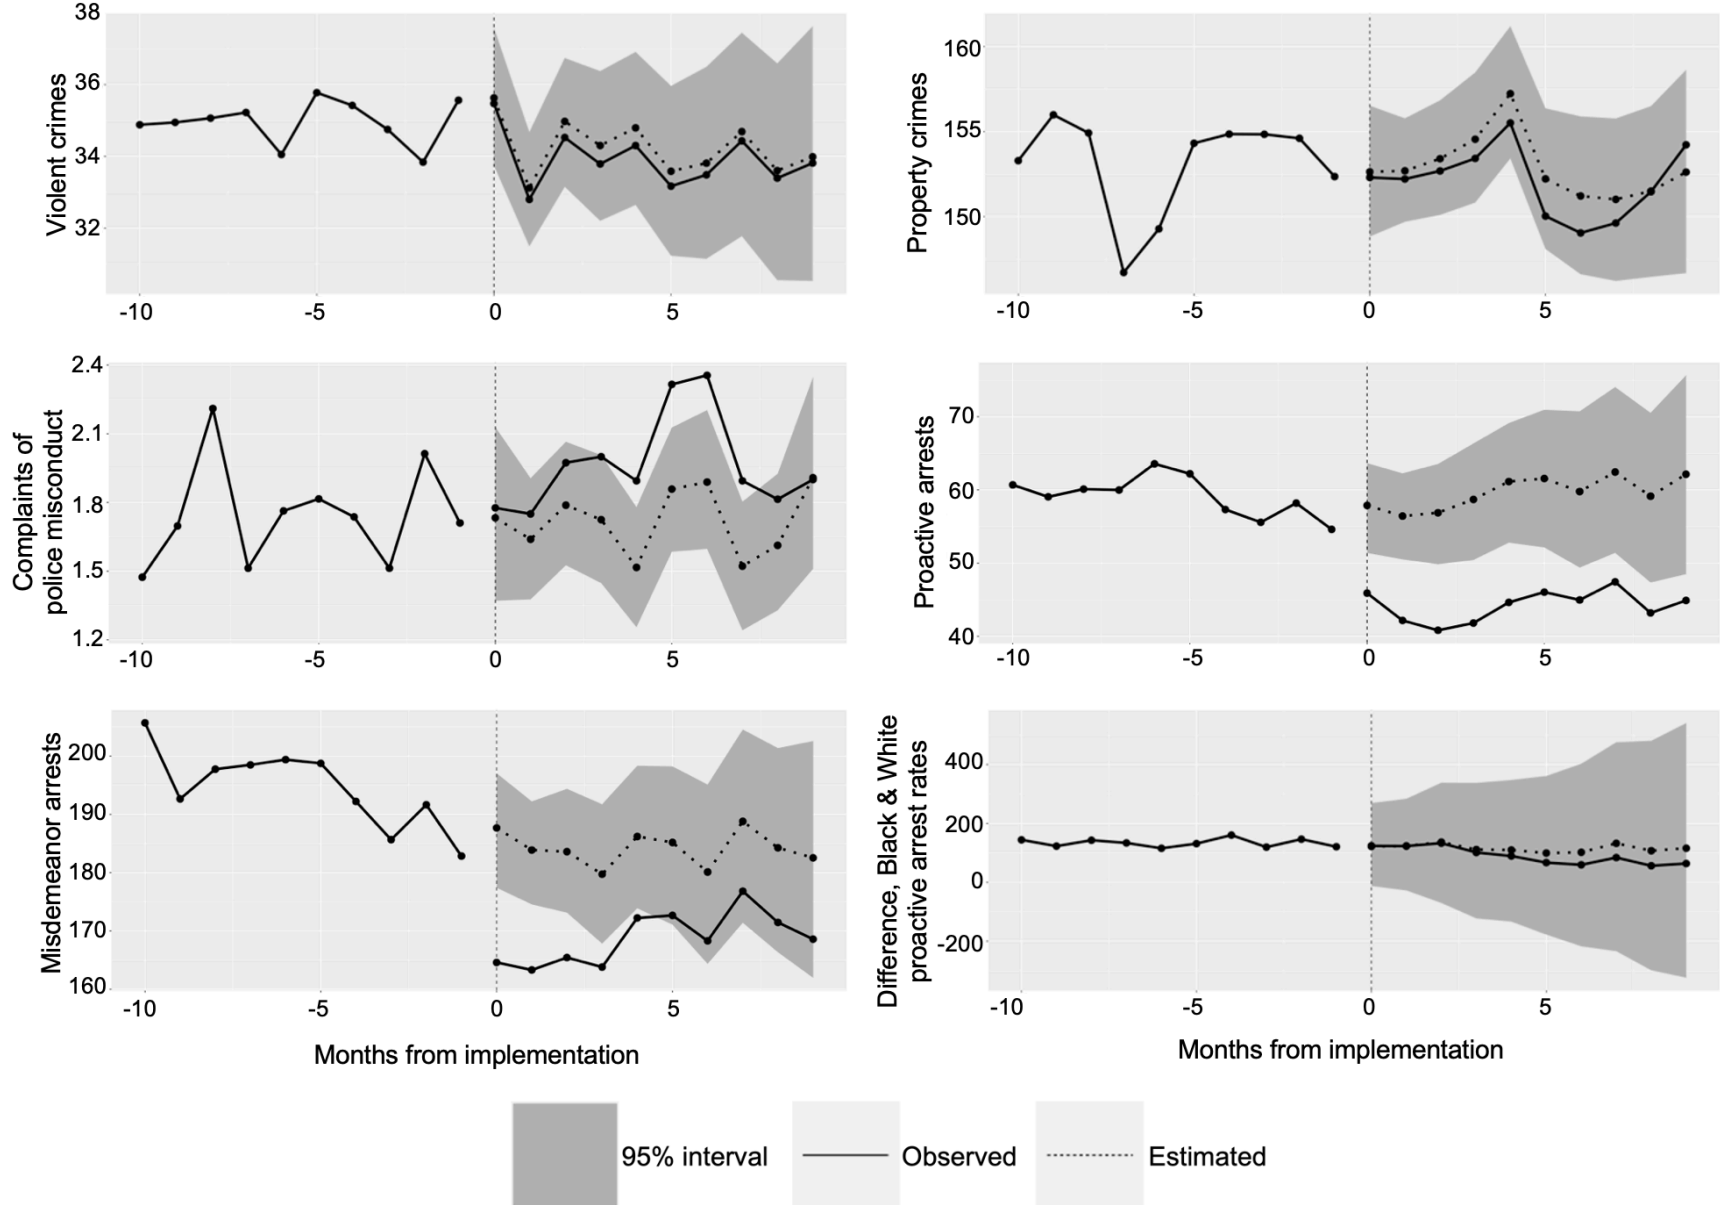


Figure 4. Significance and direction of treatment effect modifiers



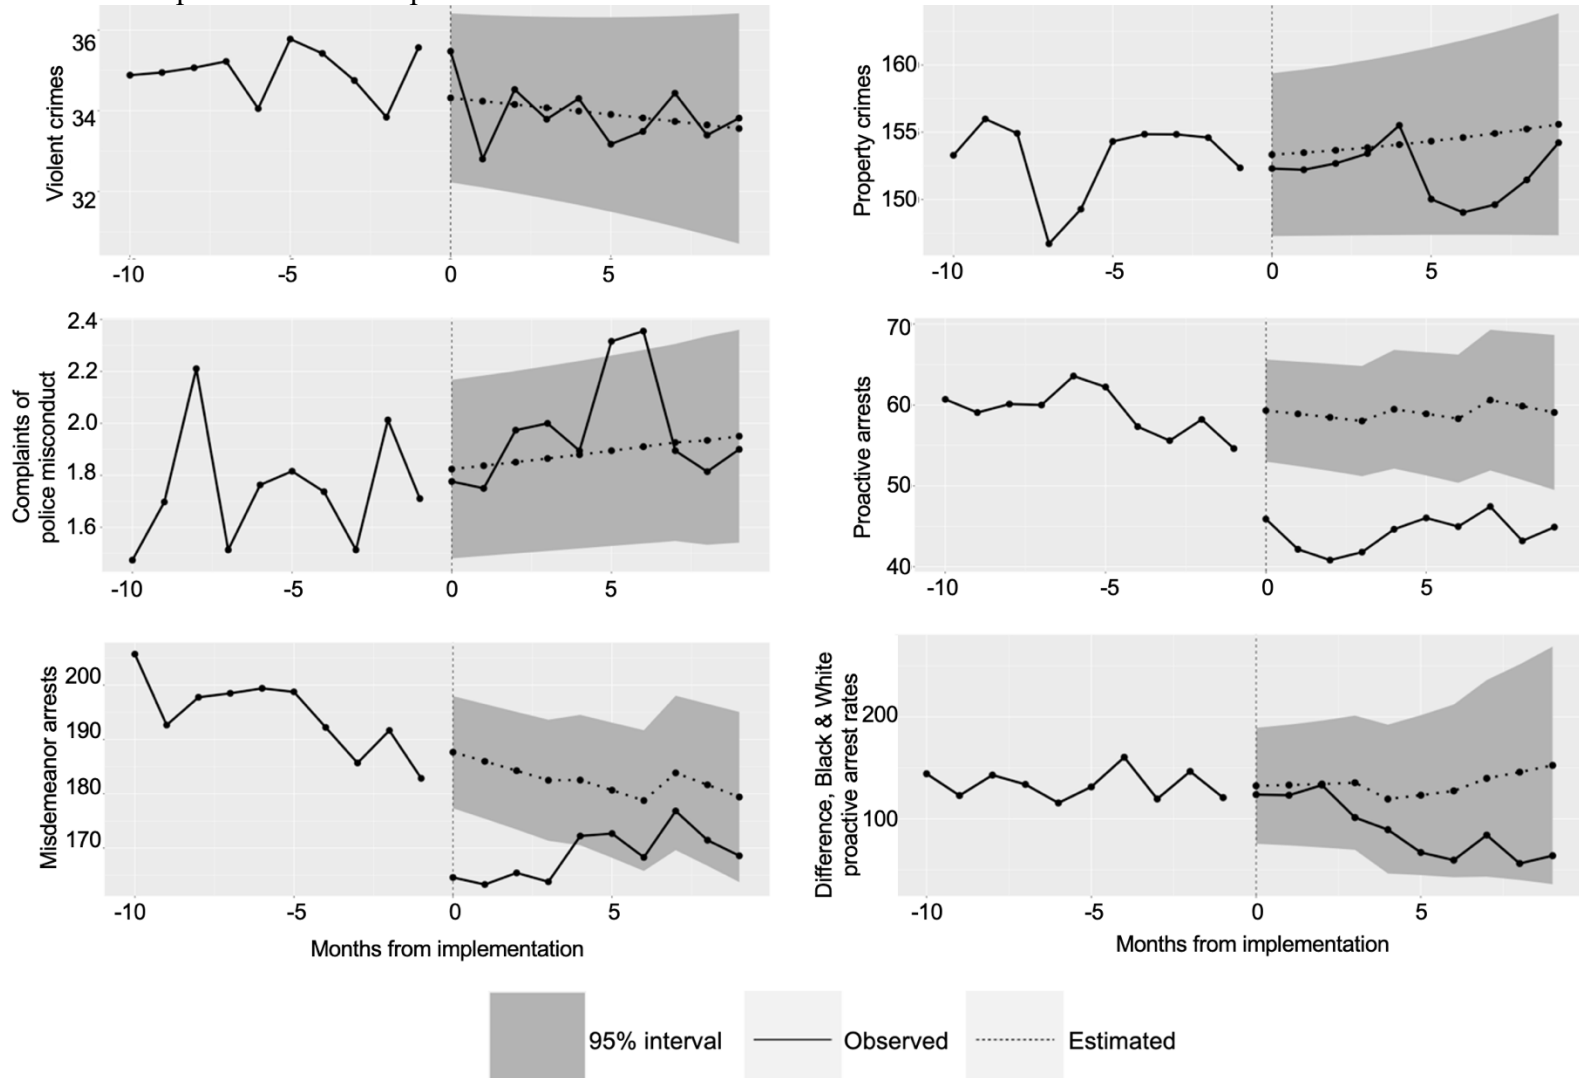


Figure 5. Treatment effect on proactive arrests in each precinct



Appendix

Appendix A. Interrupted time series replications of main models



Appendix B. Marginal effects plots, alternate outcomes

